

UTILITY PATENT APPLICATION TRANSMITTAL

(Large Entity)

(Only for new nonprovisional applications under 37 CFR 1.53(b))

Docket No.
YO999-374

Total Pages in this Submission

TO THE ASSISTANT COMMISSIONER FOR PATENTS

Box Patent Application
Washington, D.C. 20231

Transmitted herewith for filing under 35 U.S.C. 111(a) and 37 C.F.R. 1.53(b) is a new utility patent application for an invention entitled:

SYSTEM AND METHOD FOR ENFORCING COMMUNICATIONS BANDWIDTH BASED SERVICE LEVEL AGREEMENTS TO PLURALITY OF CUSTOMERS HOSTED ON A CLUSTERED WEB SERVER

and invented by:

Richard Harold Boivie, Daniel Manuel Dias, Colin George Harrison, Eric M. Levy-Abegnoli, Jean A. Lorrain, Kiyoshi Maruyama, and Pascal Pol Marie Thubert

If a CONTINUATION APPLICATION, check appropriate box and supply the requisite information:

☐ Continuation ☐ Divisional ☐ Continuation-in-part (CIP) of prior application No.: _____

Which is a:

☐ Continuation ☐ Divisional ☐ Continuation-in-part (CIP) of prior application No.: _____

Which is a:

☐ Continuation ☐ Divisional ☐ Continuation-in-part (CIP) of prior application No.: _____

Enclosed are:

Application Elements

1. ☒ Filing fee as calculated and transmitted as described below
2. ☒ Specification having 26 pages and including the following:
 - a. ☒ Descriptive Title of the Invention
 - b. ☐ Cross References to Related Applications (if applicable)
 - c. ☐ Statement Regarding Federally-sponsored Research/Development (if applicable)
 - d. ☐ Reference to Microfiche Appendix (if applicable)
 - e. ☒ Background of the Invention
 - f. ☒ Brief Summary of the Invention
 - g. ☒ Brief Description of the Drawings (if drawings filed)
 - h. ☒ Detailed Description
 - i. ☒ Claim(s) as Classified Below
 - j. ☒ Abstract of the Disclosure

UTILITY PATENT APPLICATION TRANSMITTAL
(Large Entity)

(Only for new nonprovisional applications under 37 CFR 1.53(b))

Docket No.
Y0999-374

Total Pages in this Submission

Application Elements (Continued)

3. ☒ Drawing(s) *(when necessary as prescribed by 35 USC 113)*
- a. ☐ Formal Number of Sheets _____
- b. ☒ Informal Number of Sheets 5 (Figs. 1-4B)
4. ☒ Oath or Declaration
- a. ☒ Newly executed *(original or copy)* ☐ Unexecuted
- b. ☐ Copy from a prior application (37 CFR 1.63(d)) *(for continuation/divisional application only)*
- c. ☒ With Power of Attorney ☐ Without Power of Attorney
- d. ☐ DELETION OF INVENTOR(S)
Signed statement attached deleting inventor(s) named in the prior application,
see 37 C.F.R. 1.63(d)(2) and 1.33(b).
5. ☐ Incorporation By Reference *(usable if Box 4b is checked)*
The entire disclosure of the prior application, from which a copy of the oath or declaration is supplied
under Box 4b, is considered as being part of the disclosure of the accompanying application and is hereby
incorporated by reference therein.
6. ☐ Computer Program in Microfiche *(Appendix)*
7. ☐ Nucleotide and/or Amino Acid Sequence Submission *(if applicable, all must be included)*
- a. ☐ Paper Copy
- b. ☐ Computer Readable Copy *(identical to computer copy)*
- c. ☐ Statement Verifying Identical Paper and Computer Readable Copy

Accompanying Application Parts

8. ☐ Assignment Papers *(cover sheet & document(s))*
9. ☐ 37 CFR 3.73(B) Statement *(when there is an assignee)*
10. ☐ English Translation Document *(if applicable)*
11. ☐ Information Disclosure Statement/PTO-1449 ☐ Copies of IDS Citations
12. ☐ Preliminary Amendment
13. ☒ Acknowledgment postcard
14. ☐ Certificate of Mailing
- ☐ First Class ☐ Express Mail *(Specify Label No.):* _____

UTILITY PATENT APPLICATION TRANSMITTAL
(Large Entity)

(Only for new nonprovisional applications under 37 CFR 1.53(b))

Docket No.
YO999-374

Total Pages in this Submission

Accompanying Application Parts (Continued)

15. ☐ Certified Copy of Priority Document(s) *(if foreign priority is claimed)*


16. ☐ Additional Enclosures *(please identify below):*

Fee Calculation and Transmittal

CLAIMS AS FILED

For	#Filed	#Allowed	#Extra	Rate	Fee
Total Claims	27	- 20 =	7	x \$18.00	\$126.00
Indep. Claims	8	- 3 =	5	x \$78.00	\$390.00
Multiple Dependent Claims (check if applicable) <input type="checkbox"/>					\$0.00
BASIC FEE					\$690.00
OTHER FEE (specify purpose) _____					\$0.00
TOTAL FILING FEE					\$1,206.00

- ☒ A check in the amount of **\$1,206.00** to cover the filing fee is enclosed.
- ☒ The Commissioner is hereby authorized to charge and credit Deposit Account No. **50-0481** as described below. A duplicate copy of this sheet is enclosed.
- ☐ Charge the amount of _____ as filing fee.
- ☒ Credit any overpayment.
- ☒ Charge any additional filing fees required under 37 C.F.R. 1.16 and 1.17.
- ☐ Charge the issue fee set in 37 C.F.R. 1.18 at the mailing of the Notice of Allowance, pursuant to 37 C.F.R. 1.311(b).


Signature

Dated: **February 18, 2000**

Sean M. McGinn, Esq.
Reg. No.: 34,386

CC:

Customer No.: 21254

MCGINN & GIBB, P.C.
A PROFESSIONAL LIMITED LIABILITY COMPANY
PATENTS, TRADEMARKS, COPYRIGHTS, AND INTELLECTUAL PROPERTY LAW
1701 CLARENDON BOULEVARD, SUITE 100
ARLINGTON, VIRGINIA 22209
TELEPHONE (703) 294-6699
FACSIMILE (703) 294-6696

**APPLICATION
FOR
UNITED STATES
LETTERS PATENT**

APPLICANTS: **Richard Harold Boivie, Daniel Manuel
Dias, Colin George Harrison, Eric M.
Levy-Abegnoli, Jean A. Lorrain, Kiyoshi
Maruyama, and Pascal Pol Marie
Thubert**

FOR: **SYSTEM AND METHOD FOR
ENFORCING COMMUNICATIONS
BANDWIDTH BASED SERVICE
LEVEL AGREEMENTS TO
PLURALITY OF CUSTOMERS
HOSTED ON A CLUSTERED WEB
SERVER**

DOCKET NO.: **YO999-374**

**SYSTEM AND METHOD FOR ENFORCING COMMUNICATIONS BANDWIDTH
BASED SERVICE LEVEL AGREEMENTS TO PLURALITY OF CUSTOMERS
HOSTED ON A CLUSTERED WEB SERVER**

Field of the Invention

5 The present invention relates generally to a world-wide network, and more particularly to those of Internet World Wide Web (WWW) sites of various owners hosted by a service provider using a cluster or group of servers and meeting with agreed upon service levels.

Background of the Invention

10 The Internet is the world's largest network, and it has become essential in organizations such as government, academia, small, medium and large businesses as well as to individual consumers. Many businesses have started out-sourcing their electronic business ("e-business") and electronic commerce ("e-commerce") Web sites to service providers instead of running their Web sites on their own server(s) and managing them by themselves.

15 Such a service provider needs to install a collection of servers (termed a Web Server Farm (WSF), Universal Server Farm (USF), or Web Server Cluster) which can be used by many different businesses to support their e-commerce and e-business. These business customers (e.g., the service provider's "customers") have different "capacity" requirements for their Web sites. The users of e-commerce (consumers, business partners, etc.) access the Web Server Farm by

first logging into the Internet. The Web Server Farm (WSF) is connected to the Internet via high speed communications links such as T3 and OCx links. These links are shared by all the Web sites, and all the users who are accessing the services hosted by the Web Server Farm.

When businesses out-source their e-commerce and/or e-business to a service provider, they must get some guarantee on the services they are getting from the service provider for their sites. Once the service provider makes a commitment to a customer to provide a certain "level" of service (e.g., termed a Service Level Agreement (SLA)), the provider must guarantee that level of service to that customer.

A general service level agreement (SLA) on communications link bandwidth usage for a customer can be denoted by a pair of bandwidth constraints: the minimum guaranteed bandwidth $B(i, \min)$ and the maximum bandwidth bound $B(i, \max)$ for each i -th customer.

The minimum (or min) bandwidth $B(i, \min)$ is a guaranteed bandwidth that the i -th customer will receive regardless of the bandwidth usage by other customers. The maximum (or max) bandwidth $B(i, \max)$ is an upper bound on the bandwidth that the i -th customer may receive provided that some unused bandwidth is available (e.g., bandwidth not being currently used by other customers).

Therefore, the range between $B(i, \min)$ and $B(i, \max)$ represents bandwidth provided on an "available" or "best-effort" basis to a customer, and it is not necessarily guaranteed that the customer will obtain this bandwidth.

Generally, the unit cost to use the bandwidth up to $B(i, \min)$ is less than or equal to the unit cost to use the bandwidth between $B(i, \min)$ and $B(i, \max)$. Such a unit cost assigned to one customer may differ from those assigned to other customers.

In the environment of Web site hosting, where communications link(s) between the Internet and a server farm is shared by a number of customers (i.e., traffic to/and from the customer Web sites are sharing the communications link(s)), the bandwidth management on the outbound link (i.e., the link from a server farm to the Internet) is more important than the bandwidth management on the inbound link since the amount of traffic on the outbound link is several orders of magnitude greater than that on the inbound link.

Furthermore, in most cases, the inbound traffic to the server farm is directly responsible for the outbound traffic generated by the server farm. Therefore, the service level agreements ($B(i, \min)$, $B(i, \max)$) are generally applied on the outbound link bandwidth usage.

A conventional method for leaky bucket traffic shaping uses fair queuing collision arbitration. This method uses a set of queues and virtual finishing time for scheduling Asynchronous Transfer Mode (ATM) cell traffic. The method could be implemented to regulate packet traffic to enforce the minimum bandwidth SLA, $B(i, \min)$ to each customer.

Another conventional method for shaping traffic in a packet-switched network uses a set of queues and the conformance time to shape packet traffic. The method can be used to regulate packet traffic to enforce the minimum bandwidth SLA, $B(i, \min)$ to each customer.

Some commercial products (e.g., Xedia Access Point (www.Xedia.com) and Phasecom's SpeedClass (www.speed-demon.com) can be used to regulate directly the outbound traffic to meet with the (minimum, maximum) bandwidth SLA for each customer.

However, while the conventional systems and methods such as those mentioned above, could be reasonably applied to enforce service level agreements (SLAs) on the outbound link

usage by each customer (and on customer traffic class or type), some can only support the minimum bandwidth SLA and some can support the (minimum, maximum) bandwidth SLA.

A major problem of these systems and methods is that they enforce the outbound bandwidth SLA by throttling (i.e., dropping some of) the traffic already generated by specific source (IP) addresses.

A major problem arises when some packets must be dropped because either the outbound link has been congested for a sustained period of time so that some queues are running out of space, and/or a particular customer's outbound link usage has been exceeding the maximum bandwidth SLA and its queue(s) has filled up. That is, when a queue has filled up, packets coming to that queue are randomly dropped. This means the packet-dropping affects many responses from a server farm since each response consists of many packets. This dropping of packets triggers many Transmission Control Protocol (TCP) retransmission, which leads to even further retransmission of packets for error recovery, causing thrashing and then eventually leading to a slowdown of many connections or disconnection of connections.

Another problem of dropping outbound packets is wasting server resources in generating responses which cannot be delivered to end users (i.e., servers are kept busy for generating undeliverable responses, yet they could have been used for generating responses for customers whose outbound bandwidth usage are below the minimum guarantees).

Yet another problem is, classification of outbound traffic is limited since the outbound packet does not say for which type request the packet was generated. This will limit the degree to which differentiated services could be applied in controlling bandwidth usage.

SUMMARY OF THE INVENTION

In view of the foregoing and other problems, disadvantages and drawbacks of the conventional systems and methods, an object of the present invention is a system and method to guarantee and deliver the (minimum, maximum) based communications link bandwidth Service Level Agreements (SLAs) to customers whose applications (e.g., Web sites) are hosted by a server or a server cluster.

Another object of the present invention is to prevent any single customer's traffic in the server or server cluster from monopolizing (e.g., "hogging") the entire bandwidth resource and penalizing others.

Yet another object of the present invention is to provide a feedback method and system for enforcing the outbound link bandwidth SLAs by throttling and shaping the inbound traffic to a server or a server cluster.

A further object of the present invention is to provide differentiated services to various types of traffic within a defined class of traffic by supporting multilevel selective- packet-drop criteria.

A still further object of the present invention is to serve the incoming requests by priority order and not necessarily in their order of arrival.

Yet another object of the present invention is to provide a method and system for classifying and shaping traffic using various levels from the Internet Protocol (IP) addresses to User Resource Locators (URLs).

Another object of the invention is to provide a method and system for maximizing the usage of the outbound link bandwidth and to provide fair usage to various classes of traffic.

Yet another object of the present invention is to provide a method and system for maximizing the potential service revenue in allocating unused bandwidth to various classes of traffic.

A still further object of the invention is to provide a method and system for avoiding the generation of output from a server or a server cluster if that output could not be delivered to end users.

Another object of the present invention is to provide a method and system for dropping requests and/or connections to minimize the negative business impact to customers.

In a first aspect of the present invention, a system and method are provided to guarantee and deliver the communications link bandwidth-based Service Level Agreements (SLAs) to customers whose applications (e.g., Web sites) are hosted by a server or a server farm. A method of controlling and guaranteeing a service level agreement (SLA), based on a communications outbound link bandwidth usage to a plurality of customers having electronic business activity hosted by at least one server as a server farm, includes a) monitoring the outbound communications bandwidth usage by each customer traffic to determine a level of service being provided to each customer with respect to the agreed service level agreement in each service cycle time per unit of time; and b) controlling a flow of incoming requests to each customer business activity application so as to guarantee a level of service previously agreed to the customer by queuing requests and by selectively dropping requests to the customer to guarantee the agreed service levels to the customer. The controlling step controls and guarantees

each outbound link usage based service level agreement by controlling the flow of incoming requests to the at least one server.

The Service Level Agreements preferably are in the form of (minimum, maximum) bandwidth bounds where the minimum bandwidth represents the guaranteed bandwidth while the maximum bandwidth represents the upper bound to the as-available use bandwidth. A customer may have an SLA for each type traffic. The system enforces the outbound link bandwidth SLAs by throttling and shaping inbound traffic to a server farm.

Preferably, incoming traffic is classified into various class/type queues. Each queue has multiple water-levels to trigger selective-dropping of packets prior to dropping any new arriving packets. Traffic in queues is scheduled and shipped to a server farm using a real-time scheduling algorithm that support and enforce the (minimum, maximum) bandwidth-based SLA to each customer traffic. The system uses the information on the outbound link bandwidth usage by various class traffic as a primary means to throttle and shape inbound traffic to a server farm. The scheduling algorithm looks for an opportunity to increase the revenue to a server farm operator when allocating bandwidth resources beyond the minimum bandwidth SLAs.

With the unique and unobvious structure and method of the invention, the (minimum, maximum) based communications link bandwidth SLAs are guaranteed and delivered to customers whose applications are hosted by a server or a server cluster. Further, single customer traffic (a class traffic belonging to a single customer) in the server or server cluster are prevented from monopolizing an entire bandwidth resource at the expense of other customers. Along these lines, the usage of the outbound link bandwidth is maximized and yet fair usage is provided to

various classes of traffic. Additionally, feedback is provided for enforcing the outbound link bandwidth SLAs by throttling and shaping the inbound traffic to a server or a server cluster.

Additionally, with the present invention, differentiated services are provided to various types of traffic within a defined class of traffic by supporting multilevel selective- packet-drop criteria. Further, incoming requests are served by priority order and not necessarily in their order of arrival. Additionally, traffic is classified and shaped using various levels such as from the Internet Protocol (IP) addresses to User Resource Locators (URLs). The potential service revenue is maximized in allocating unused bandwidth to various classes of traffic. Further, the output generation from a server or a server cluster is avoided if that output cannot be delivered to end users. Thus, server resources are conserved, and unnecessary output (packets) are not generated and subsequently dropped. Further, requests and/or connections are dropped to minimize the negative business impact to customers.

BRIEF DESCRIPTION OF THE DRAWINGS

The foregoing and other objects, aspects and advantages will be better understood from the following detailed description of a preferred embodiment of the invention with reference to the drawings, in which:

Figure 1 illustrates an overall system and environment of the present invention;

Figure 2 illustrates a communications bandwidth manager 110 according to the present invention for managing and monitoring a communications bandwidth on an output data path;

Figure 3 illustrates the components of the Communications Bandwidth Manager 110 including a traffic estimator 340, a scheduler 320, and a queuing system 360;

Figure 4A illustrates a flowchart of an inbound traffic scheduling method 400; and

Figure 4B illustrates a method of computing $b(i)$, the expected bandwidth usage per a
5 request that is admitted to a server farm.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS OF THE PRESENT INVENTION

Referring now to the drawings, and more particularly to Figures 1-7, an embodiment of the present invention will be described herein below.

10 Figure 1 shows the overall system 100 and environment for the present invention. As shown in Figure 1, clients 130 (e.g., Client_1, Client_2, Client_3, ..., Client_n) on a world-wide network (e.g., the Internet) 120, typically running Web browsers, make requests for web data located on a cluster of Web servers 100.

The requests arrive at a Communications Bandwidth Manager (CBM) 110. The CBM
15 110 presents a single IP address to the Web server cluster. The CBM 110 may queue incoming requests in a set of queues (as described in further detail below) in order to satisfy service level agreements (also described below). The CBM selects a request from the queues, selects one of the servers 101 (e.g., Server_1, Server_2, Server_3, etc.) to service the request, and sends the request to that server.

20 The selected server receives the Web request, services it, and sends the response directly back to the client 130 (e.g., Client_3) along path 160. The return path 160 between the servers

100 and the Internet 120, is shared by the cluster of servers. This path 160 carries the bulk of the Web server bandwidth load.

The present invention is directed at controlling the allocation of this link bandwidth among multiple customer Web sites hosted on the Web server cluster 100, by controlling
5 incoming Web requests at the CBM 110.

Figure 2 illustrates the main concept of the present invention embodied in the Communications Bandwidth Manager (CBM) 110. The CBM 110 receives incoming (i.e., inbound) Web requests 240 from Web clients 130, on input link 241.

The input requests 240 are queued in a set of queues 250 (e.g., in order of arrival time)

10 The requests in a queue 250 are serviced first-in, first out (FIFO) within each particular class (type) of requests. While Figure 2 only shows a single queue 250, there exists queues for each traffic class for each customer. That is, a customer may have a single queue or multiple queues. If there are multiple queues, there are queues for each traffic type (e.g., class, priority, etc.). Thus, within each of the classes, the requests are serviced in a FIFO manner. Once, the request
15 arrives in FIFO, the request is placed into a predetermined queue based on the customer, different customer's Web site, server, etc. Hence, within each particular queue, the FIFO processing is preserved but not across different queues.

In addition to Web input requests, the input requests 240 include packets associated with the ongoing Web requests, such as acknowledgment packets.

20 The acknowledgment packets contain information on the quantity of outbound data that is being acknowledged, as specified in the conventional protocols, such as the specification of the

TCP/IP (Transmission Control Protocol / Internet Protocol). This information may be used to estimate the volume of data that was sent out on the outgoing path 160 from the server 101.

As illustrated in Figure 2, the outgoing path 160 may be monitored, in place of monitoring acknowledgment packets (e.g., generated by the server and indicating, for example, the number of data bytes delivered/received by a customer).

The path 160 carries the outbound traffic (i.e., responses to requests). Hence, monitoring the output on the outgoing data path 160 can be monitored and it can be observed how much output is being generated. If the amount of bandwidth being used by a particular customer exceeds the number of bandwidth per the SLA, then feedback can be generated to reduce the number of inbound requests being accepted.

Thus, the present invention provides a feedback loop, so as to provide monitored data 260 indicating the bandwidth on output path 160 being used by each of the Web sites hosted on the Web server cluster 100.

This monitored data 260 may be obtained directly by monitoring the output path 160 (e.g., monitoring the flow and the available bandwidth) at each of the Web servers in the cluster.

Based on the monitored data 260 (or the data contained in acknowledgment packets), configuration information (e.g., the connectivity between the CBM and servers) and accumulated information (e.g., server performance data) based on operation of the cluster, the CBM 110 selects from the queued incoming requests 240 (e.g., based on the exemplary scheduling processes shown in Figures 4A and 5A which describe the criteria for selecting certain requests over others), selects a Web server node 101 (e.g., based on a determination made

by a conventional load balancer or the like which selects one node over another) to service the request, and forwards the request 280 to the selected server 101.

Based on the load of the outgoing path 160, some incoming requests may accumulate in queues 250. Thus, if traffic is heavy (e.g., above a predetermined threshold), some or all of the

5 queues 250 may overflow, or some incoming requests may get delayed for a time period beyond a threshold. In these cases, the CBM 110 selects one or more incoming requests and discards them. Hence, once the queue is backed up, then there may be a delay of the incoming requests and a discarding of requests. Upon the queue becoming full, the incoming requests coming to a particular queue may be discarded in order of arrival. It is noted that the scheduler determines
10 whether to admit (or not admit) a request. Typically, requests in a queue will not be dropped, but only those request newly-arriving will be dropped.

Alternatively, for these “to-be discarded” requests, the CBM may either directly, or via a selected server, return a message to the Web client 130 that the server complex is overloaded.

There may be a dedicated server to return such a message (e.g., not to process requests, but only
15 to return such “Server Busy” messages). Once, for example, a browser receives such a message, the browser can proceed to a next stage. Otherwise, the browser must wait until a “time out” message is received. After the “time out” the browser would move to a next process.

For example, in an Internet browser, when “CNN” is clicked, the “CNN” server will bring up a template, and within the template there are instructions to the Internet browser to
20 “click” on various fields so that the template can be filled with content (e.g., images, text, etc.). Such “clicks” by the browser are made transparent to the user, to obtain the plurality of information. When the browser cannot obtain something, then the browser must retransmit the

request. However, to avoid this, the dedicated server described above could be employed, thereby freeing the browser for other tasks.

Figure 3 shows the components of the CBM 110. Incoming Web requests on input link 241 are placed in the queuing system 360. Monitored data 260 on the usage of the outgoing data path 160 arrive, and are gathered, at the traffic estimator 340. The traffic estimator may also gather the output load information from acknowledgment packets that arrive at the queuing system 360, via link 310.

The traffic estimator compiles the load information, and inputs the composed information to a scheduler component 350 over link 320. Based on the data from the traffic estimator, service level agreement (SLA) information (e.g., provided by the server(s)), and other configuration information (e.g., the connectivity between the CBM and servers), the scheduler 350 selects a request in the queuing system 360, and determines the server node 101 to service the request, and sends the selected request to the server node over link 280.

The scheduler 350 also determines if requests or other associated packets in the queuing system 360 are to be discarded, or whether a direct response indicating an overloaded condition is to be sent to the client Web browser.

Exemplary Scheduling Method

Figure 4A shows the scheduling method/algorithm 400 for inbound requests to a server farm after the requests have been queued in various queues 360.

The scheduling method 400 starts at step 401. The rest of steps (e.g., 402 through 409) will be repeated in a preset service cycle time. Step 402 resets control variables such as $b(i)$, B ,

$B(i)$ and $R(i)$ for all class i . Here, $b(i)$ represents the expected outbound link bandwidth usage in a service cycle time had a request in the class i been admitted to the server farm, B represents the current available bandwidth space that is set to B_{max} at the beginning of each service cycle time (e.g., B_{max} is the maximum usable bandwidth available on the outbound link in a service cycle time), $B(i)$ represents the expected bandwidth usage had $R(i)$ number of requests been dispatched to a server farm and $R(i)$ represents the number of requests scheduled (and dispatched) thus far within the cycle time.

Step 403 determines whether or not a given service cycle time has expired. If it has expired (e.g., “YES”), then the method loops back to step 402. If the given service cycle time has not yet expired (e.g., “NO”), then step 404 checks whether or not any bandwidth is left for use. If no space is available (i.e., B is less-than-or-equal-to 0, denoted by $B \leq 0$), step 405 is executed, and then goes back to step 402.

If some bandwidth space is available (i.e., $B > 0$, B greater than zero), then the algorithm executes 406 and checks if there is any queue i such that a request is in that queue and its planned bandwidth usage is less than ($<LT>$) its minimum bandwidth SLA, $B(i, \min)$.

Step 406 guarantees the minimum bandwidth SLA. If the condition of the step 406 is “YES”, then the algorithm schedules (and dispatches) one request from queue i and updates control variables; $B(i) = B(i) + b(i)$, $R(i) = R(i) + 1$, and $B = B - b(i)$, and then goes back to step 403.

If the condition in step 406 is “NO” (i.e., as of this moment, every class has received the minimum bandwidth SLA service), then step 407 is executed, in which the algorithm determines

whether there is class i such that there is a request in that queue and its bandwidth usage is less than its maximum bandwidth SLA, $B(i, \max)$.

If the condition in step 407 is "NO", then the method loops back to step 403. If the condition 407 is "YES", then step 408 is executed and class i is selected. The selection of a class
5 is based on any given criteria predetermined by the designer. Examples of these criteria are priorities, weights assigned to classes, which class pays more than others, which class should receive preferential treatment over others, etc.

Once a class is selected in step 408, step 409 is executed to update the control variables, and then loops back to step 403. Steps 406, 407 and 408 provide the best effort services defined
10 by the range between the minimum bandwidth SLA, $B(i, \min)$ and the maximum bandwidth SLA, $B(i, \max)$.

That is, for purposes of the invention, "best effort services" means that for each customer there is an SLA. Within the SLA, the customer is typically guaranteed a minimum bandwidth, or rate (bytes/bits per second) generated and delivered by their server. Thus, there is a minimum
15 guaranteed level of service (e.g., 1 MBits per second guaranteed). There also typically is a range of additional between the minimum guarantee and a maximum bandwidth which the customer may be offered (e.g., 2 Mbits per second for a total of 3 Mbits per second), to service additional users coming into the customer's Web server and thus the Web server may generate a lot more than 1 Mbit per second of traffic. The range between 1 Mbit per second (minimum bandwidth)
20 and 3 Mbit per second (maximum) is termed the "best efforts" range (also the "as available" range). When the customer starts using the bandwidth in this range, the customer is charged for

such usage. This bandwidth beyond the minimum is not guaranteed and the customer may not get it, depending upon other customer's usage and SLAs.

Figure 4B describes how $b(i)$ in step 402 of Figure 4A can be computed. $b(i)$ is derived by dividing $Ba(i)$ by $Ra(i)$ where $Ba(i)$ represents the averaged/smoothed bandwidth usage that is derived from the measured or monitored bandwidth usage $Bm(i)$ in the immediately preceding cycle time, and $Ra(i)$ represents the averaged/smoothed number of requests.

Figure 4B shows one way of smoothing by using a constant "a" where a is a value between zero and one. The averaging/smoothing is used to absorb "spikes" (e.g., mitigate the effect of spikes on the waveform) to derive statistically meaningful values. There are many other conventional means to obtain the averaged or smoothed value as are known by one of ordinary skill in the art .

The scheduling method described above and shown in Figure 4A attempts to utilize the currently- available bandwidth as much as possible.

The method 400 can be modified easily to be applied to control and regulate the outbound link bandwidth usage directly to meet the (min, max) SLAs, similar approach used by the conventional products. Such a modification is related to the computation of $b(i)$ in Step 402 in Figure 4A and in Figure 4B. Instead of "computing" $b(i)$, the value $b(i)$ can be "obtained" from the packet that is being scheduled and dispatched. (i.e., each packet carries its length information, and thus this packet length information can be used as $b(i)$ in stead of computing $b(i)$ as described in Figure 4b) This is possible since the method is applied directly to the traffic on the outbound link from the server farm to the Internet.

The CBM described in this invention can be implemented by hardware or software or both. For example, the CBM may be implemented by operating a computer, as embodied by a digital data processing apparatus, to execute a sequence of machine-readable instructions. These instructions may reside in various types of programmable signal-bearing media.

5 Thus, this aspect of the present invention is directed to a programmed product, including signal-bearing media tangibly embodying a program of machine-readable instructions executable by a digital data processor to perform the above method.

10 Hence, in addition to the hardware and process environment described above, a different aspect of the invention includes a computer-implemented method according to the present invention, as described above, for enforcing communications bandwidth based on service level agreements to a plurality of customer hosted on a clustered web server, as described above. As an example, this method may be implemented in the particular hardware environment by executing a sequence of machine-readable instructions in the signal-bearing media.

15 With the present invention, a closed-loop control is provided whereby the outbound traffic (e.g., responses) provides the control of the incoming traffic (e.g., requests).

 While the invention has been described in terms of preferred embodiments, those skilled in the art will recognize that the invention can be practiced with modification within the spirit and scope of the appended claims.

CLAIMS

Having thus described our invention, what we claim as new and desire to secure by Letters Patent is as follows:

1. A method of controlling and guaranteeing a service level agreement (SLA) based on a
5 communications outbound link bandwidth usage to a plurality of customers having electronic business activity hosted by at least one server as a server farm, said method comprising:

a) monitoring said outbound communications bandwidth usage by each customer traffic to determine a level of service being provided to each customer with respect to the agreed service level agreement in each service cycle time per unit of time;

10 b) controlling a flow of incoming requests to each customer business activity application so as to guarantee a level of service previously agreed to said customer by queuing and by selectively dropping requests to said customer to guarantee the agreed service levels to said customer,

wherein said controlling controls and guarantees each outbound link usage based
15 service level agreement by controlling the flow of incoming requests to a server or group of servers.

2. The method as set forth in Claim 1, wherein said communications resource for service level agreements comprises an outbound communications link from said at least one server to said

network which sends responses from said server farm.

3. The method as set forth in Claim 1, wherein said communications resource customer traffic generated comprises a server farm having a plurality of servers, which is controlled by throttling the traffic to said server farm.

5 4. The method as set forth in Claim 1, wherein said communications resource customer traffic generated comprises a server farm, having at least one server, which is controlled by throttling the traffic to said server farm.

5. A method of controlling and guaranteeing the service level agreement (SLA) based on the communications outbound link bandwidth usage to a plurality of customers whose e-business
10 and e-commerce are hosted by a server or a set of servers as a server farm, said method comprising:

a) monitoring said outbound communications bandwidth usage by each customer traffic to determine the level of service being provided to each customer with respect to the agreed service level agreement in each "service cycle time" per unit of time;

15 b) controlling the flow of incoming requests to each customer eBusiness/eCommerce application so as to guarantee a level of service previously agreed to said customer by queuing requests to said customer and by selectively dropping requests to said customer to guarantee the agreed service levels to said customer.

6. A method of guaranteeing a service level agreement (SLA) based on a communications outbound link bandwidth usage to a plurality of customers having electronic business activity hosted by at least one server, said method comprising:

a) monitoring said outbound communications bandwidth usage by each customer traffic

5 to determine a level of service being provided to each customer with respect to the agreed service level agreement in each service cycle time per unit of time; and

b) controlling a flow of incoming requests to each said customer business activity application so as to guarantee a level of service previously agreed to said customer by queuing requests to said customer and by selectively dropping requests to said customer to guarantee the
10 agreed service levels to said customer,

wherein said controlling controls and guarantees each outbound link usage based service level agreement by controlling the flow of incoming requests to the at least one server.

7. A method of regulating inbound requests on a world-wide network, comprising:

monitoring an amount of inbound traffic requests on a link of said world-wide network

15 for a plurality of customers; and

regulating an output generated based on said amount of inbound traffic requests monitored in order to meet a service level agreement for said plurality of customers.

8. A system for controlling and guaranteeing a service level agreement (SLA) based on a communications outbound link bandwidth usage to a plurality of customers having electronic

20 business activity hosted by at least one server as a server farm, said system comprising:

a) means for monitoring said outbound communications bandwidth usage by each customer traffic to determine a level of service being provided to each customer with respect to the agreed service level agreement in each service cycle time per unit of time; and

b) means for controlling a flow of incoming requests to each customer business activity application so as to guarantee a level of service previously agreed to said customer by queuing requests to said customer and by selectively dropping requests to said customer to guarantee the agreed service levels to said customer.

9. A communications system, comprising:

a worldwide network for communicating with a plurality of customers;

a manager, operatively coupled to said worldwide network, for controlling and guaranteeing a service level agreement (SLA) based on a communications outbound link bandwidth usage to said plurality of customers; and

at least one server functioning as a server farm, operatively coupled to said manager, said plurality of customers having electronic business activity hosted by said at least one server as said server farm.

10. The system according to claim 9, wherein said worldwide network comprises the Internet.

11. The system according to claim 10, wherein said customers output requests on said Internet for web data located on a cluster of servers.

12. The system according to claim 9, wherein said manager comprises a Communications Bandwidth Manager (CBM) presenting a single address to the at least one server cluster.

13. The system according to claim 12, wherein said manager includes a set of queues for queuing incoming requests.

5 14. The system according to claim 13, wherein said manager selects a request from the queues, selects one of the at least one server to service the request, and sends the request to that server.

15. The system according to claim 14, wherein said selected server receives the request, services said request, and sends a response directly back to said customer along a data output path,

10 wherein a portion of the path between said at least one server and said worldwide network is shared by a cluster of servers.

16. The system according to claim 15, wherein said manager controls an allocation of said outgoing data path among multiple customer sites hosted on the server cluster, by controlling incoming requests at the manager.

17. The system according to claim 9, wherein said manager comprises:

15 an input link for receiving incoming requests from customers; and
a set of queues for queuing the incoming requests.

18. The system according to claim 17, wherein said incoming requests are queued in order of arrival time and are serviced first-in, first out (FIFO) within each predetermined traffic class of requests.

19. The system according to claim 18, wherein said queues are provided for each traffic class for
5 each customer.

20. The system according to claim 19, wherein the incoming requests include acknowledgment packets contain information on the quantity of outbound data that is being acknowledged, so as to estimate a volume of data that was output on an outgoing path from the server.

21. The system according to claim 19, wherein said manager monitors an outgoing data path to
10 determine a number of data units delivered to a customer.

22. The system according to claim 21, wherein said manager, based on monitoring of said outgoing data path, determines if an amount of bandwidth being used by a customer exceeds a bandwidth amount per a service level agreement, such that feedback is generated to reduce the number of inbound requests being accepted.

23. The system according to claim 19, wherein said manager further comprises a traffic
15 estimator for gathering monitored data on the usage of an output data path,

said traffic estimator also gathering output load information from said acknowledgment packets that arrive at the set of queues.

24. The system according to claim 23, further comprising a scheduler for receiving compiled load information from said traffic estimator,

5 wherein, based on the data from the traffic estimator, and service level agreement (SLA) information provided by an operator, said scheduler selects from requests in the queues, and for a selected request determines a server node to service the request, and sends the selected request to the server node.

25. The system according to claim 24, wherein said scheduler determines if requests or
10 associated packets in the queues are to be discarded, or whether a direct response indicating an overloaded condition is to be sent to a customer.

26. A signal-bearing medium tangibly embodying a program of machine-readable instructions executable by a digital processing apparatus to perform a method of controlling and guaranteeing a service level agreement (SLA) based on a communications outbound link bandwidth usage to a
15 plurality of customers having electronic business activity hosted by at least one server as a server farm, said method comprising:

a) monitoring said outbound communications bandwidth usage by each customer traffic to determine a level of service being provided to each customer with respect to the agreed service level agreement in each service cycle time per unit of time; and

b) controlling a flow of incoming requests to each customer business activity application so as to guarantee a level of service previously agreed to said customer by queuing requests to said customer and by selectively dropping requests to said customer to guarantee the agreed service levels to said customer.

5 27. A signal-bearing medium tangibly embodying a program of machine-readable instructions executable by a digital processing apparatus to perform a method of regulating inbound requests on a world-wide network, said method comprising:

monitoring an amount of inbound traffic requests on a link of said world-wide network for a plurality of customers; and

10 regulating an output generated based on said amount of inbound traffic requests monitored in order to meet a service level agreement for said plurality of customers.

**SYSTEM AND METHOD FOR ENFORCING COMMUNICATIONS BANDWIDTH
BASED SERVICE LEVEL AGREEMENTS TO PLURALITY OF CUSTOMERS
HOSTED ON A CLUSTERED WEB SERVER**

ABSTRACT OF THE DISCLOSURE

5 A method and system for controlling and guaranteeing a service level agreement (SLA)
based on a communications outbound link bandwidth usage to a plurality of customers having
electronic business activity hosted by at least one server as a server farm, includes monitoring the
outbound communications bandwidth usage by each customer traffic to determine a level of
service being provided to each customer with respect to the agreed service level agreement in
10 each service cycle time per unit of time. The flow of incoming requests to each customer
business activity application is controlled so as to guarantee a level of service previously agreed
to the customer by queuing requests to the customer and by selectively dropping requests to the
customer to guarantee the agreed service levels to the customer. The controlling process controls
and guarantees each outbound link usage based service level agreement by controlling the flow
15 of incoming requests to the at least one server.

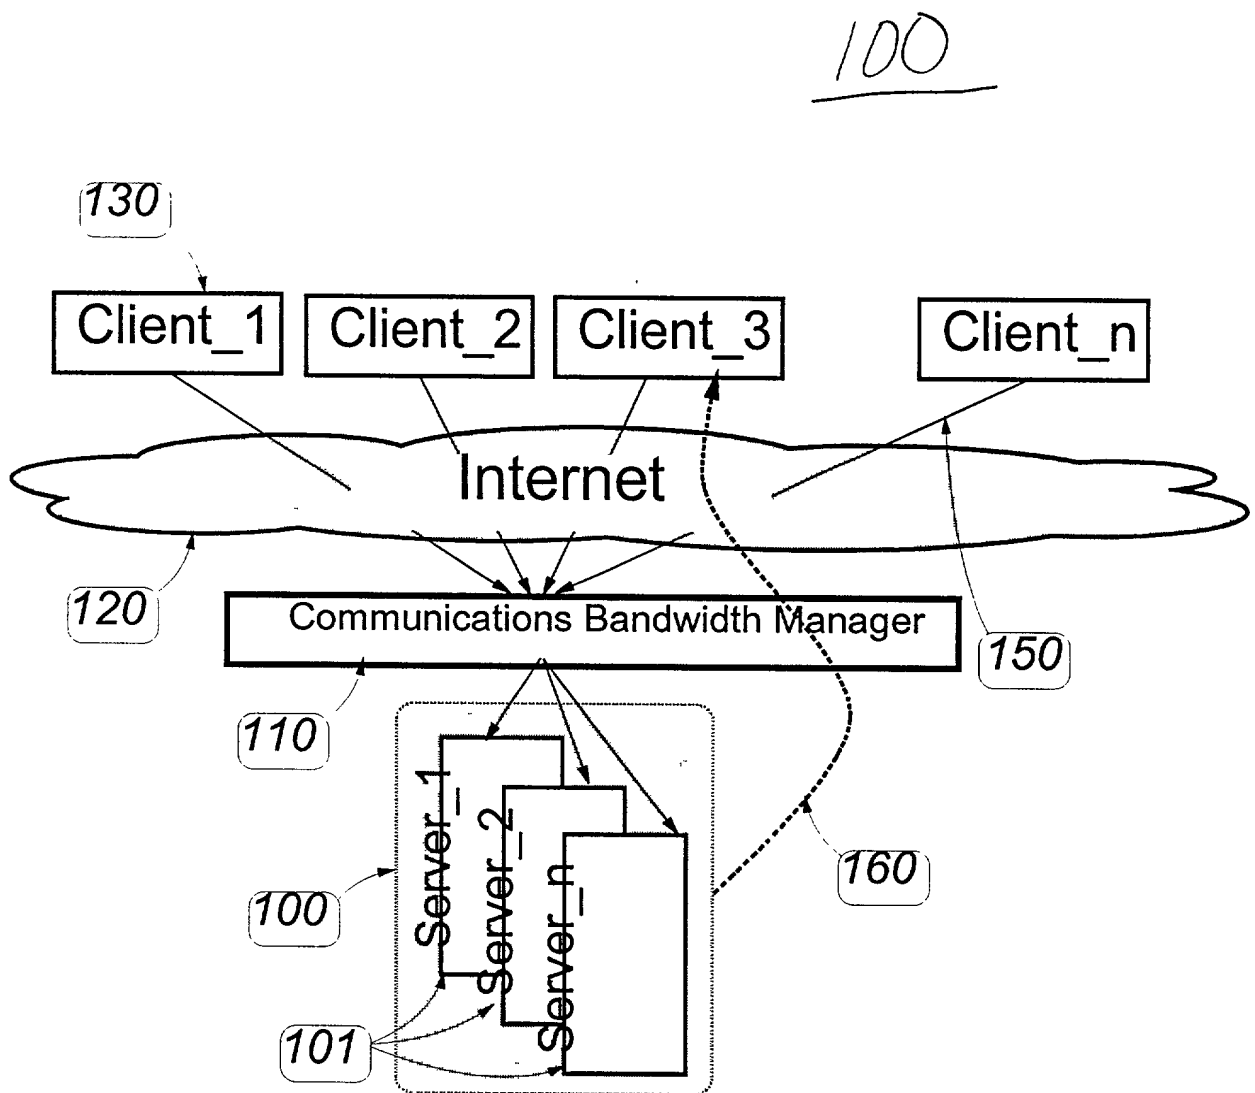


Figure 1




Figure 2


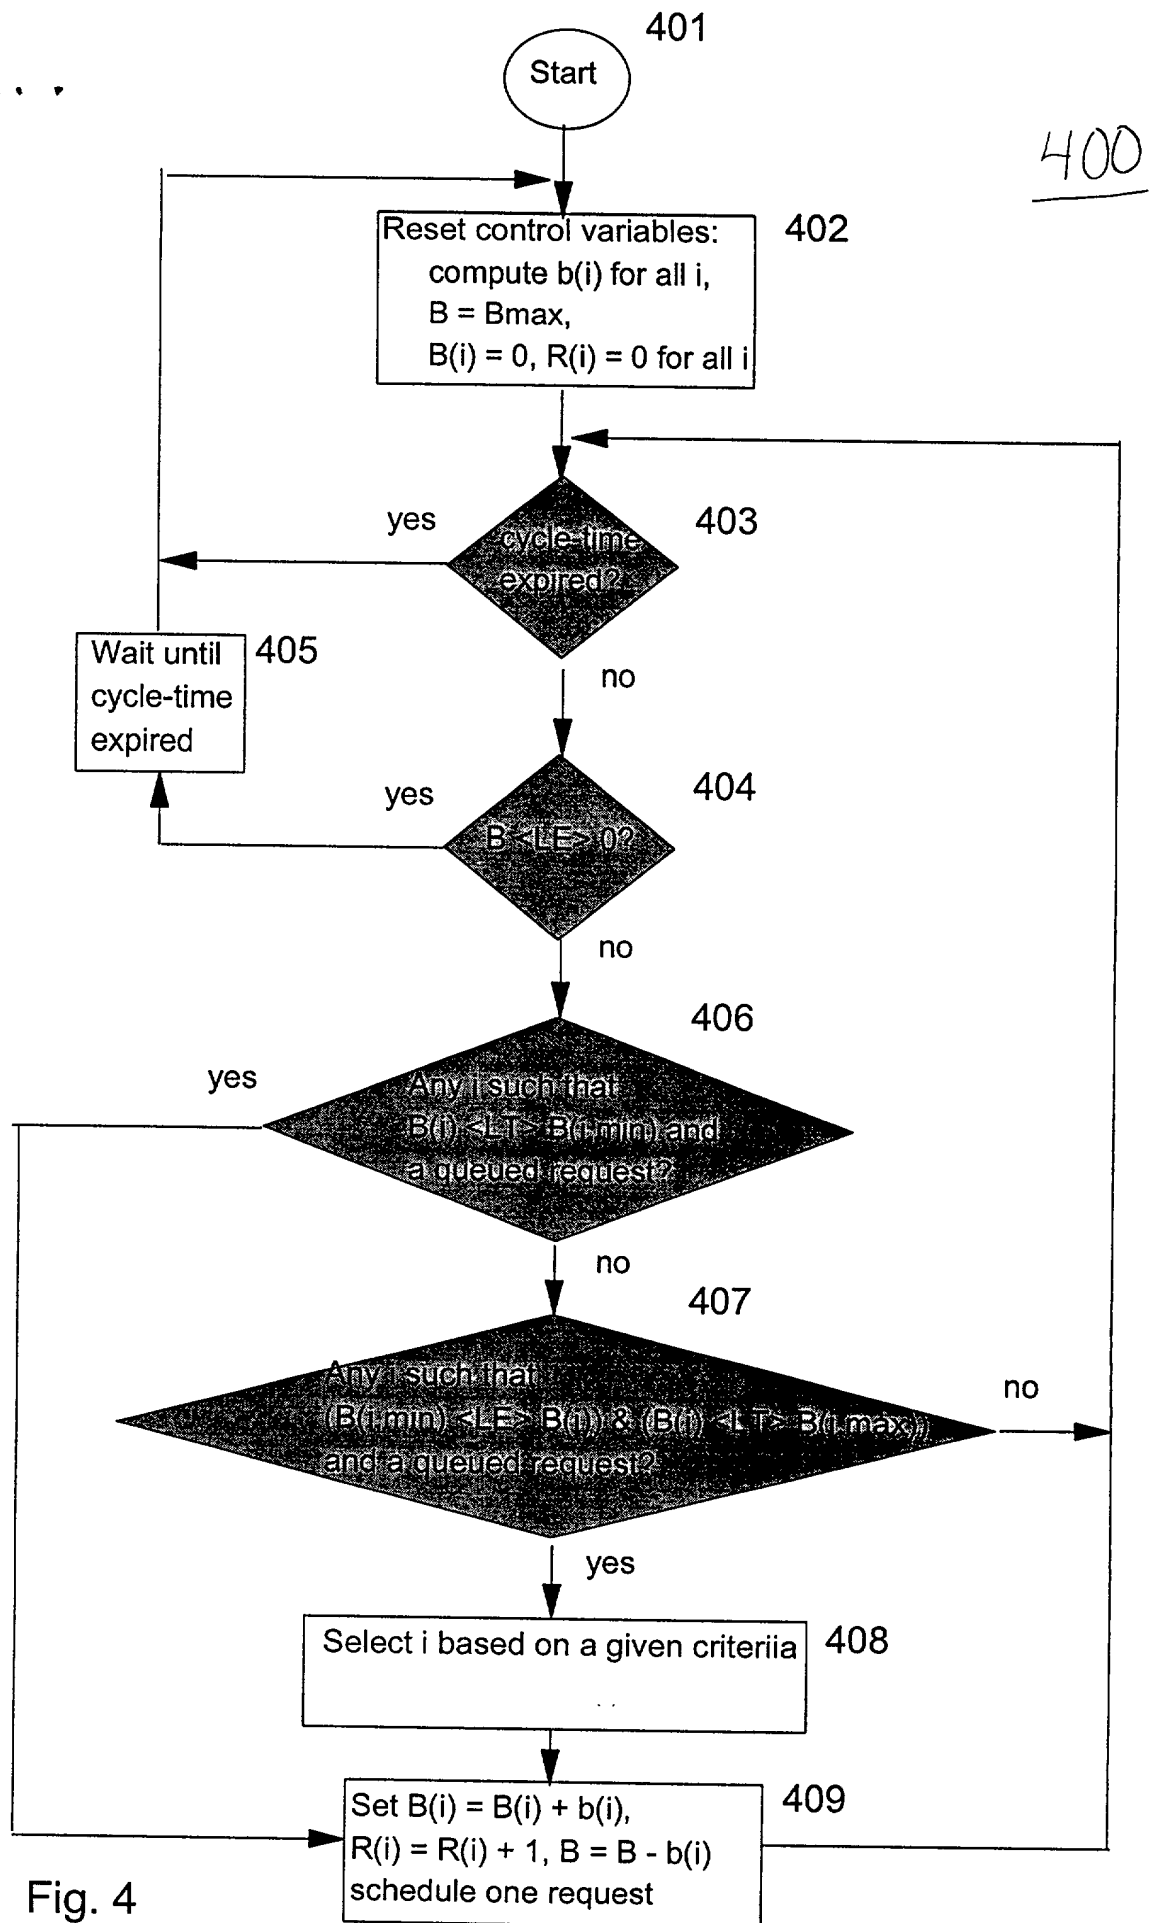


Figure 3.



For all i

$$Ra(i) = (1 - a)Ra(i) + aR(i),$$

$$Ba(i) = (1 - a)Ba(i) + aBm(i),$$

$$b(i) = Ba(i) / Ra(i).$$

where

$Ra(i)$ represents averaged/smoothed number of requests

$Ba(i)$ represents averaged/smoothed bandwidth usage

$Bm(i)$ represents the measured/monitored bandwidth usage

all in a service cycle time, and

$b(i)$ represents the expected bandwidth usage by a one request

Fig. 4B

IBM Docket No.: YO999-374

DECLARATION AND POWER OF ATTORNEY

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name: I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled: **SYSTEM AND METHOD FOR ENFORCING COMMUNICATIONS BANDWIDTH BASED SERVICE LEVEL AGREEMENTS TO PLURALITY OF CUSTOMERS HOSTED ON A CLUSTERED WEB SERVER**

the specification of which:
(check one)

☒ is attached hereto.

☐ was filed on _____, as Application Serial No. _____ and was amended on _____.

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose information which is material to the patentability of this application in accordance with Title 37, Code of Federal Regulations, § 1.56.

I hereby claim foreign priority benefits under Title 35, United States Code, § 119 of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed:

Prior Foreign Application(s):

Number	Country	Day/Month/Year	Priority Claimed
--------	---------	----------------	------------------

I hereby claim the benefit under Title 35, United States Code, § 120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code, § 112, I acknowledge the duty to disclose material information as defined in Title 37, Code of Federal Regulations, § 1.56 which occurred between the filing date of the prior application and the national or PCT international filing date of this application.

Prior U.S. Applications:

Serial No.	Filing Date	Status
------------	-------------	--------

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

As a named inventor, I hereby appoint the following attorneys and/or agents to prosecute this application and transact all business in the Patent and Trademark Office connected therewith: We hereby appoint Masny Schechter, Registration No. 31,722, Tony J. Iardi, Registration No. 29,936, Christopher A. Hughes, Registration No. 26,914, Edward A. Pennington, Registration No. 32,588, John E. Hoel, Registration No. 26,279, Joseph C. Redmond, Jr., Registration No. 18,753, Douglas W. Cameron, Registration No. 31,596, Louis P. Herzberg, Registration No. 41,500, Kevin M. Jordan, Registration No. 40,277, Stephen C. Kaufman, Registration No. 29,551, Daniel P. Morris, Registration No. 32,053, Louis J. Perocchio, Registration No. 33,206, Jay P. Strulzini, Registration No. 36,266, David M. Shoff, Registration No. 39,635, Paul J. Overstreet, Registration No. 37,411 and Robert M. Trepp, Registration No. 25,933, to prosecute this application and transact all business in the United States Patent and Trademark Office connected therewith.

Send all correspondence to: Sean M. McGinn, McGinn & Gibb, P.C., 1701 Clarendon Boulevard, Suite 100, Arlington, Virginia 22209. Customer No. 21254

Telephone calls should be directed to Sean M. McGinn, McGinn & Gibb, P.C. at (703) 294-6699.

(1) Inventor: Richard Harold Borvie

Signature: Richard Harold BorvieDate: 2/14/2000

Residence: 194 Outlets Farm Road, Monroe, CT 06468

Citizenship: United States of America

Post Office Address: Same as Residence

IBM DocId: No.: Y0999-374

- (2) Inventor: Daniel Manuel Dias
Signature: Daniel Manuel Dias Date: 2/14/2000
Residence: 3380 Sunay Court, Mohagan Lake, Yorktown, New York 10598
Citizenship: United States of America
Post Office Address: Same as Residence
- (3) Inventor: Colin George Harrison
Signature: Colin George Harrison Date: 02/14/00
Residence: 16 Cove Road, Brookfield, CT 06804
Citizenship: United Kingdom
Post Office Address: Same as Residence
- (4) Inventor: Eric M. Levy-Abeignoli
Signature: Eric M. Levy-Abeignoli Date: 02/15/2000
Residence: Residence Alexandra, 67 Avenue Chemin de la Lanterne, 06200 Nice, France
Citizenship: France
Post Office Address: Same as Residence
- (5) Inventor: Jean A. Lorrain
Signature: Jean A. Lorrain Date: 02/16/2000
Residence: 2 Avenue Marie-Alexandre, 06140 Vence, France
Citizenship: France
Post Office Address: Same as Residence
- (6) Inventor: Kiyoshi Maruyama
Signature: Kiyoshi Maruyama Date: 2-14-2000
Residence: 7 Green Lane, Chappaqua, New York 10516
Citizenship: Japan
Post Office Address: Same as Residence
- (7) Inventor: Pascal Pol Marie Thubert
Signature: Pascal Pol Thubert Date: 2-16-2000
Residence: Les Jardins d'Elise, 60 Av Des Poilus, 06140 Vence, France
Citizenship: France
Post Office Address: Same as Residence

** TOTAL PAGE.07 **